arXiv:1403.6602v1 [cs.DS] 26 Mar 2014

Pivot Sampling in Java 7's Dual-Pivot Quicksort

Exploiting Asymmetries in Yaroslavskiy's Partitioning Scheme

Markus E. Nebel[1,2†] and Sebastian Wild[1]

[1] *Computer Science Department, University of Kaiserslautern*
[2] *Department of Mathematics and Computer Science, University of Southern Denmark*

Abstract: The new dual-pivot Quicksort by Vladimir Yaroslavskiy — used in Oracle's Java runtime library since version 7 — features intriguing asymmetries in its behavior. They were shown to cause a basic variant of this algorithm to use less comparisons than classic single-pivot Quicksort implementations. In this paper, we extend the analysis to the case where the two pivots are chosen as fixed order statistics of a random sample and give the precise leading term of the average number of comparisons, swaps and executed Java Bytecode instructions. It turns out that — unlike for classic Quicksort, where it is optimal to choose the pivot as median of the sample — the asymmetries in Yaroslavskiy's algorithm render pivots with a systematic skew more efficient than the symmetric choice. Moreover, the optimal skew heavily depends on the employed cost measure; most strikingly, abstract costs like the number of swaps and comparisons yield a very different result than counting Java Bytecode instructions, which can be assumed most closely related to actual running time.

Keywords: Quicksort, dual-pivot, Yaroslavskiy's partitioning method, median of three, average case analysis

1 Introduction

Quicksort is one of the most efficient comparison-based sorting algorithms and is thus widely used in practice, for example in the sort implementations of the C++ standard library and Oracle's Java runtime library. Almost all practical implementations are based on the highly tuned version of Bentley and McIlroy (1993), often equipped with the strategy of Musser (1997) to avoid quadratic worst case behavior. The Java runtime environment was no exception to this — up to version 6. With version 7 released in 2009, however, Oracle broke with this tradition and replaced its tried and tested implementation with a new dual-pivot Quicksort variant proposed by Vladimir Yaroslavskiy.

The decision was based on extensive running time experiments that clearly favored the new algorithm. This was particularly remarkable as earlier analyzed dual-pivot Quicksort variants had not shown any potential for performance gains over classic single-pivot Quicksort (Sedgewick, 1975; Hennequin, 1991). Unlike those variants, however, Yaroslavskiy's partitioning method makes clever use of asymmetries in the algorithm and thereby needs less comparisons: asymptotically $1.9\,n\ln n$ vs. $2\,n\ln n$ comparisons on

[†] The order of authors follows the Hardy-Littlewood rule, *i. e.*, it is alphabetical by last name.

1365–8050 © 2005 Discrete Mathematics and Theoretical Computer Science (DMTCS), Nancy, France

average (Wild and Nebel, 2012). These savings are opposed by a much larger number of swaps, so that the overall competition remained open. A comparison of Java implementations of the two Quicksort variants showed that Yaroslavskiy's method actually executes *more* Java Bytecode instructions on average (Wild et al., 2013b).

All of the above results ignore a very effective strategy for improving the basic Quicksort algorithm: For decades, practical implementations choose their pivots as *median of a random sample* of the input to be more efficient (both in terms of average performance and in making worst cases less likely). Oracle's Java 7 implementation also employs this optimization: It chooses its two pivots as the *tertiles of five* sample elements. This equidistant choice is a plausible generalization, since selecting the pivot as median is known to be optimal for classic Quicksort (Sedgewick, 1975; Martínez and Roura, 2001).

However, the classic partitioning methods treat elements smaller and larger than the pivot in symmetric ways — unlike Yaroslavskiy's partitioning algorithm: Depending on how elements relate to the two pivots, one of *five* different execution paths is taken in the partitioning loop, and these have highly different costs (Wild et al., 2013c)! How often each of these five paths is taken thus depends on the ranks of the two pivots, which we can push in a certain direction by selecting *other* order statistics of a sample than the tertiles. The partitioning costs alone are then minimized if the cheapest execution path is taken all the time. This however leads to very unbalanced distributions of sizes for the recursive calls, such that a *trade-off* between partitioning costs and balance of subproblem sizes results. **It will be the purpose of this paper to identify the optimal way to sample pivots by means of a precise analysis of the resulting overall costs,** and to validate earlier empirical findings that way.

Related work. Single-pivot Quicksort with pivot sampling has been intensively studied over the last decades (Emden, 1970; Sedgewick, 1975, 1977; Hennequin, 1991; Martínez and Roura, 2001; Neininger, 2001; Chern and Hwang, 2001; Durand, 2003). We heavily profit from the mathematical foundations laid by these authors. There are scenarios where, even for the symmetric, classic Quicksort, a skewed pivot can yield benefits over median of k (Martínez and Roura, 2001; Kaligosi and Sanders, 2006). An important difference to Yaroslavskiy's algorithm is, however, that the situation remains symmetric: A relative pivot rank $\alpha < \frac{1}{2}$ has the same effect as one with rank $1 - \alpha$.

Aumüller and Dietzfelbinger (2013) establish a lower bound of asymptotically $1.8\, n \ln n$ comparisons for any partitioning method for dual-pivot Quicksort. Kushagra et al. (2014) show that Yaroslavskiy's algorithm needs fewer scans over the array, and thus less cache misses, than classic Quicksort; this is a possible explanation for better running times. Both references focus on Quicksort without pivot sampling.

Outline. After some general notation, Section 3 introduces the subject of study. Section 4 collects the main analytical results of this paper, whose proof is divided into Sections 5, 6 and 7. Arguments in the main text are kept concise, but the interested reader is provided with details in the appendix. The algorithmic consequences of our analysis are discussed in Section 8. Section 9 concludes the paper.

2 Notation and Preliminaries

We write vectors in bold font, for example $\mathbf{t} = (t_1, t_2, t_3)$. For concise notation, we use expressions like $\mathbf{t} + 1$ to mean *element-wise* application, *i.e.*, $\mathbf{t} + 1 = (t_1 + 1, t_2 + 1, t_3 + 1)$. By $\mathrm{Dir}(\boldsymbol{\alpha})$, we denote a random variable with *Dirichlet distribution* and shape parameter $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_d) \in \mathbb{R}^d_{>0}$. Likewise for parameters $n \in \mathbb{N}$ and $\mathbf{p} = (p_1, \dots, p_d) \in [0,1]^d$ with $p_1 + \cdots + p_d = 1$, we write $\mathrm{Mult}(n, \mathbf{p})$ for a random variable with *multinomial distribution* with n trials. $\mathrm{HypG}(k, r, n)$ is a random variable with

hypergeometric distribution, *i. e.*, the number of red balls when drawing k times without replacement from an urn of $n \in \mathbb{N}$ balls, r of which are red, (where $k, r \in \{1, \dots, n\}$). Finally, $\mathcal{U}(a, b)$ is a random variable uniformly distributed in the interval (a, b), and $\mathrm{B}(p)$ is a Bernoulli variable with probability p to be 1. We use "$\stackrel{\mathcal{D}}{=}$" to denote equality in distribution.

As usual for the average case analysis of sorting algorithms, we assume the *random permutation model*, *i. e.*, all elements are different and every ordering of them is equally likely. The input is given as array $\mathtt{A}$ of length n and we denote the initial entries of $\mathtt{A}$ by $U_1, \dots, U_n$. We further assume that $U_1, \dots, U_n$ are i. i. d. uniformly $\mathcal{U}(0, 1)$ distributed; as their ordering forms a random permutation (Mahmoud, 2000), this assumption is without loss of generality.

3 Generalized Yaroslavskiy Quicksort

In this section, we review Yaroslavskiy's partitioning method and combine it with the pivot sampling optimization to obtain what we call the *Generalized Yaroslavskiy Quicksort* algorithm. We leave some parts of the algorithm unspecified here, but give a full-detail implementation in the appendix. The reason is that *preservation of randomness* is somewhat tricky to achieve in presence of pivot sampling, but vital for precise analysis. The casual reader might content him- or herself with our promise that everything turns out alright in the end; the interested reader is invited to follow our discussion of this issue in Appendix A.

3.1 Generalized Pivot Sampling

Our pivot selection process is declaratively specified as follows, where $\mathbf{t} = (t_1, t_2, t_3) \in \mathbb{N}^3$ is a fixed parameter: Choose a random sample $\mathbf{V} = (V_1, \dots, V_k)$ of size $k = k(\mathbf{t}) := t_1 + t_2 + t_3 + 2$ from the elements and denote by $(V_{(1)}, \dots, V_{(k)})$ the *sorted*[(i)] sample, *i. e.*, $V_{(1)} \le V_{(2)} \le \cdots \le V_{(k)}$. Then choose the two pivots $P := V_{(t_1+1)}$ and $Q := V_{(t_1+t_2+2)}$ such that they divide the sorted sample into three regions of respective sizes t_1, t_2 and t_3:

$$\underbrace{V_{(1)} \dots V_{(t_1)}}_{t_1 \text{ elements}} \le \underbrace{V_{(t_1+1)}}_{=P} \le \underbrace{V_{(t_1+2)} \dots V_{(t_1+t_2+1)}}_{t_2 \text{ elements}} \le \underbrace{V_{(t_1+t_2+2)}}_{=Q} \le \underbrace{V_{(t_1+t_2+3)} \dots V_{(k)}}_{t_3 \text{ elements}} . \quad (3.1)$$

Note that by definition, P is the small(er) pivot and Q is the large(r) one. We refer to the $k - 2$ elements of the sample that are not chosen as pivot as *"sampled-out"*; P and Q are the chosen *pivots*. All other elements — those which have not been part of the sample — are referred to as *ordinary* elements. Pivots and ordinary elements together form the set of *partitioning elements*, (because we exclude sampled-out elements from partitioning).

3.2 Yaroslavskiy's Dual Partitioning Method

In bird's-eye view, Yaroslavskiy's partitioning method consists of two indices, k and g, that start at the left resp. right end of $\mathtt{A}$ and scan the array until they meet. Elements left of k are smaller or equal than Q, elements right of g are larger. Additionally, a third index ℓ lags behind k and separates elements smaller than P from those between both pivots. Graphically speaking, the invariant of the algorithm is as follows:

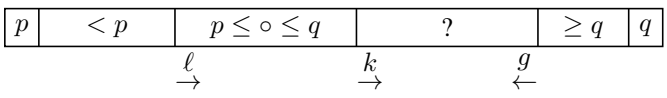

(i) In case of equal elements any possible ordering will do. However in this paper, we assume distinct elements.

We write $\mathcal{K}$ and $\mathcal{G}$ for the sets of all indices that k resp. g attain in the course of the partitioning process. Moreover, we call an element *small*, *medium*, or *large* if it is smaller than P, between P and Q, or larger than Q, respectively. The following properties of the algorithm are needed for the analysis, (see Wild and Nebel (2012); Wild et al. (2013b) for details):

(Y1) Elements U_i, $i \in \mathcal{K}$, are first compared with P. Only if U_i is not small, it is also compared to Q.

(Y2) Elements U_i, $i \in \mathcal{G}$, are first compared with Q. If they are not large, they are also compared to P.

(Y3) Every small element eventually causes one swap to put it behind ℓ.

(Y4) The large elements located in $\mathcal{K}$ and the non-large elements in $\mathcal{G}$ are always swapped in pairs.

For the analysis of the number of comparisons and swaps we will thus need to count, *e. g.*, the small elements U_i with $i \in \mathcal{K}$; we abbreviate their number by "$s@\mathcal{K}$". Similarly, $l@\mathcal{K}$ and $s@\mathcal{G}$ denote the number of large elements in k's range resp. the number of small elements in g's range.

When partitioning is finished, k and g have met and thus ℓ and g divide the array into three ranges, containing the small, medium resp. large (ordinary) elements, which are then sorted recursively. For subarrays with at most w elements, we switch to Insertionsort, (where w is constant and at least k). The resulting algorithm, Generalized Yaroslavskiy Quicksort with pivot sampling parameter $\mathbf{t} = (t_1, t_2, t_3)$ and Insertionsort threshold w, is henceforth called $Y^w_{\mathbf{t}}$.

4 Results

For $\mathbf{t} \in \mathbb{N}^3$ and $\mathcal{H}_n$ the nth harmonic number, we define the *discrete entropy* $\mathrm{H}(\mathbf{t})$ of $\mathbf{t}$ as

$$\mathrm{H}(\mathbf{t}) \;=\; \sum_{l=1}^{3} \frac{t_l+1}{k+1}(\mathcal{H}_{k+1} - \mathcal{H}_{t_l+1})\,. \tag{4.1}$$

The name is justified by the following connection between $\mathrm{H}(\mathbf{t})$ and the *entropy function* H^* of information theory: For the sake of analysis, let $k \to \infty$, such that ratios t_l/k converge to constants τ_l. Then

$$\mathrm{H}(\mathbf{t}) \;\sim\; -\sum_{l=1}^{3} \tau_l\big(\ln(t_l+1) - \ln(k+1)\big) \;\sim\; -\sum_{l=1}^{3} \tau_l \ln(\tau_l) \;=:\; \mathrm{H}^*(\boldsymbol{\tau})\,. \tag{4.2}$$

The first step follows from the asymptotic equivalence $\mathcal{H}_n \sim \ln(n)$ as $n \to \infty$. (4.2) shows that for large $\mathbf{t}$, the maximum of $\mathrm{H}(\mathbf{t})$ is attained for $\tau_1 = \tau_2 = \tau_3 = \frac{1}{3}$. Now we state our main result:

Theorem 4.1 (Main theorem): *Generalized Yaroslavskiy Quicksort with pivot sampling parameter* $\mathbf{t} = (t_1, t_2, t_3)$ *performs on average* $C_n \sim \frac{a_C}{\mathrm{H}(\mathbf{t})}\, n \ln n$ *comparisons and* $S_n \sim \frac{a_S}{\mathrm{H}(\mathbf{t})}\, n \ln n$ *swaps to sort a random permutation of* n *elements, where*

$$a_C \;=\; 1 + \frac{t_2+1}{k+1} + \frac{(2t_1+t_2+3)(t_3+1)}{(k+1)(k+2)} \quad \textit{and} \quad a_S \;=\; \frac{t_1+1}{k+1} + \frac{(t_1+t_2+2)(t_3+1)}{(k+1)(k+2)}\,. \tag{4.3}$$

Moreover, if the partitioning loop is implemented as in Appendix C of (Wild et al., 2013b), it executes on average $BC_n \sim \frac{a_{BC}}{\mathrm{H}(\mathbf{t})}\, n \ln n$ *Java Bytecode instructions to sort a random permutation of size* n *with*

$$a_{BC} \;=\; 10 \;+\; 13\frac{t_1+1}{k+1} \;+\; 5\frac{t_2+1}{k+1} \;+\; 11\frac{(t_1+t_2+2)(t_3+1)}{(k+1)(k+2)} \;+\; \frac{(t_1+1)(t_1+t_2+3)}{(k+1)(k+2)}\,. \tag{4.4}$$

The following sections are devoted to the proof of Theorem 4.1. Section 5 sets up a recurrence of costs and characterizes the distribution of costs of one partitioning step. The expected values of the latter are computed in Section 6. Finally, Section 7 provides a generic solution to the recurrence of the expected costs; in combination with the expected partitioning costs, this concludes our proof.

5 Distributional Analysis

5.1 Recurrence Equations of Costs

Let us denote by C_n the *costs* of $Y_{\mathbf{t}}^w$ on a random permutation of size n — where different "cost measures", like the number of comparisons, will take the place of C_n later. C_n is a non-negative *random* variable whose distribution depends on n. The total costs decompose into those for the first partitioning step plus the costs for recursively solving subproblems. As Yaroslavskiy's partitioning method preserves randomness (see Appendix A), we can express the total costs C_n recursively in terms of the same cost function with smaller arguments: For sizes J_1, J_2 and J_3 of the three subproblems, the costs of corresponding recursive calls are distributed like C_{J_1}, C_{J_2} and C_{J_3}, and conditioned on $\mathbf{J} = (J_1, J_2, J_3)$, these random variables are independent. Note, however, that the subproblem sizes are themselves random and inter-dependent. Denoting by T_n the costs of the first partitioning step, we obtain the following *distributional recurrence* for the family $(C_n)_{n\in\mathbb{N}}$ of random variables:

$$C_n \overset{\mathcal{D}}{=} \begin{cases} T_n + C_{J_1} + C'_{J_2} + C''_{J_3}, & \text{for } n > w; \\ W_n, & \text{for } n \le w. \end{cases} \tag{5.1}$$

Here W_n denotes the cost of Insertionsorting a random permutation of size n. $(C'_j)_{j\in\mathbb{N}}$ and $(C''_j)_{j\in\mathbb{N}}$ are independent copies of $(C_j)_{j\in\mathbb{N}}$, *i.e.*, for all j, the variables C_j, C'_j and C''_j are identically distributed and for all $\mathbf{j} \in \mathbb{N}^3$, C_{j_1}, C'_{j_2} and C''_{j_3} are totally independent, and they are also independent of T_n. We call T_n the *toll function* of the recurrence, as it quantifies the "toll" we have to pay for unfolding the recurrence once. Different cost measures only differ in the toll functions, such that we can treat them all in a uniform fashion by studying (5.1). Taking expectations on both sides, we find a recurrence equation for the *expected* costs $\mathbb{E}[C_n]$:

$$\mathbb{E}[C_n] = \begin{cases} \mathbb{E}[T_n] + \sum_{\substack{\mathbf{j}=(j_1,j_2,j_3) \\ j_1+j_2+j_3=n-2}} \mathbb{P}(\mathbf{J}=\mathbf{j})\big(\mathbb{E}[C_{j_1}] + \mathbb{E}[C_{j_2}] + \mathbb{E}[C_{j_3}]\big), & \text{for } n > w; \\ \mathbb{E}[W_n], & \text{for } n \le w. \end{cases} \tag{5.2}$$

A simple combinatorial argument gives access to $\mathbb{P}(\mathbf{J} = \mathbf{j})$: Of the $\binom{n}{k}$ different size k samples of n elements, those contribute to the probability of $\{\mathbf{J} = \mathbf{j}\}$, in which exactly t_1 of the sample elements are chosen from the overall j_1 small elements; and likewise t_2 of the j_2 medium elements and t_3 of the j_3 large ones are contained in the sample. We thus have $\mathbb{P}(\mathbf{J} = \mathbf{j}) = \binom{j_1}{t_1}\binom{j_2}{t_2}\binom{j_3}{t_3}\big/\binom{n}{k}$.

5.2 Distribution of Partitioning Costs

Let us denote by I_1, I_2 and I_3 the number of small, medium and large elements among the ordinary elements, (*i. e.*, $I_1 + I_2 + I_3 = n - k$) — or equivalently stated, $\mathbf{I} = (I_1, I_2, I_3)$ is the (vector of) sizes of the three partitions (excluding sampled-out elements). Moreover, we define the indicator variable $\delta = \mathbb{1}_{\{U_\chi > Q\}}$ to account for an idiosyncrasy of Yaroslavskiy's algorithm (see the proof of Lemma 5.1), where χ is the point where indices k and g first meet. As we will see, we can characterize the distribution of partitioning costs *conditional* on $\mathbf{I}$, *i. e.*, when considering $\mathbf{I}$ *fixed*.

5.2.1 Comparisons

For constant size samples, only the comparisons during the partitioning process contribute to the linearithmic leading term of the asymptotic average costs, as the number of partitioning steps remains linear. We can therefore ignore comparisons needed for sorting the sample. As w is constant, the same is true for subproblems of size at most w that are sorted with Insertionsort. It remains to count the comparisons during the first partitioning step.

Lemma 5.1: *Conditional on the partition sizes* $\mathbf{I}$*, the number of comparisons* $T_C = T_C(n)$ *in the first partitioning step of* $Y_{\mathbf{t}}^w$ *on a random permutation of size* $n > w$ *fulfills*

$$T_C(n) \;=\; (n-k) \;+\; I_2 \;+\; l@\mathcal{K} \;+\; s@\mathcal{G} \;+\; 2\delta \tag{5.3}$$

$$\overset{\mathcal{D}}{=}\; (n-k) \;+\; I_2 \;+\; \mathrm{HypG}(I_1+I_2, I_3, n-k) \;+\; \mathrm{HypG}(I_3, I_1, n-k) \;+\; 3\mathrm{B}\big(\tfrac{I_3}{n-k}\big)\,. \tag{5.4}$$

Proof: Every ordinary element is compared to at least one of the pivots, which makes $n - k$ comparisons. Additionally, for all medium elements, the second comparison is inevitably needed to recognize them as "medium", and there are I_2 such elements. Large elements only cause a second comparison if they are first compared with P, which happens if and only if they are located in k's range, see (Y1). We abbreviated the (random) number of large elements in $\mathcal{K}$ as $l@\mathcal{K}$. Similarly, $s@\mathcal{G}$ counts the second comparison for all small elements found in g's range, see (Y2).

The last summand 2δ accounts for a technicality in Yaroslavskiy's algorithm. If U_χ, the element where k and g meet, is large, then index k overshoots g by one, which causes two additional (superfluous) comparisons with this element. $\delta = \mathbb{1}_{\{U_\chi > Q\}}$ is the indicator variable of this event. This proves (5.3).

For the equality in distribution, recall that I_1, I_2 and I_3 are the number of small, medium and large elements, respectively. Then we need the cardinalities of $\mathcal{K}$ and $\mathcal{G}$. Since the elements right of g after partitioning are exactly all large elements, we have $|\mathcal{G}| = I_3$ and $|\mathcal{K}| = I_1 + I_2 + \delta$; (again, δ accounts for the overshoot, see Wild et al. (2013b) for detailed arguments). The distribution of $s@\mathcal{G}$, conditional on $\mathbf{I}$, is now given by the following urn model: We put all $n - k$ ordinary elements in an urn and draw their positions in $\mathtt{A}$. I_1 of the elements are colored red (namely the small ones), the rest is black (non-small). Now we draw the $|\mathcal{G}| = I_3$ elements in g's range from the urn without replacement. Then $s@\mathcal{G}$ is exactly the number of red (small) elements drawn and thus $s@\mathcal{G} \overset{\mathcal{D}}{=} \mathrm{HypG}(I_3, I_1, n-k)$.

The arguments for $l@\mathcal{K}$ are similar, however the additional δ in $|\mathcal{K}|$ needs special care. As shown in the proof of Lemma 3.7 of Wild et al. (2013b), the additional element in k's range for the case $\delta = 1$ is U_χ, which then is large by definition of δ. It thus simply contributes as additional summand: $l@\mathcal{K} \overset{\mathcal{D}}{=} \mathrm{HypG}(I_1+I_2, I_3, n-k)+\delta$. Finally, the distribution of δ is Bernoulli $\mathrm{B}\big(\frac{I_3}{n-k}\big)$, since conditional on $\mathbf{I}$, the probability of an ordinary element to be large is $I_3/(n-k)$. □

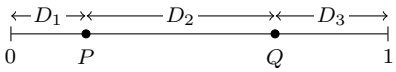

Figure 1: Graphical representation of the relation between $\mathbf{D}$ and the pivot values P and Q on the unit interval.

5.2.2 Swaps

As for comparisons, only the swaps in the partitioning step contribute to the leading term asymptotics.

Lemma 5.2: *Conditional on the partition sizes* $\mathbf{I}$*, the number of swaps* $T_S = T_S(n)$ *in the first partitioning step of* $Y_{\mathbf{t}}^w$ *on a random permutation of size* $n > w$ *fulfills*

$$T_S(n) \;=\; I_1 \,+\, l@\mathcal{K} \;\stackrel{\mathcal{D}}{=}\; I_1 \,+\, \mathrm{HypG}(I_1 + I_2, I_3, n-k) \,+\, \mathrm{B}\big(\tfrac{I_3}{n-k}\big)\,. \tag{5.5}$$

Proof: No matter where a small element is located initially, it will eventually incur one swap that puts it at its final place (for this partitioning step) to the left of ℓ, see (Y3); this gives a contribution of I_1. The remaining swaps come from the "crossing pointer" scheme, where k stops on the first large and g on the first non-large element, which are then exchanged in one swap (Y4). For their contribution, it thus suffices to count the large elements in k's range, that is $l@\mathcal{K}$. The distribution of $l@\mathcal{K}$ has already been discussed in the proof of Lemma 5.1. □

5.2.3 Bytecode Instructions

A closer investigation of the partitioning method reveals the number of executions for every single Bytecode instruction in the algorithm. Details are omitted here; the analysis is very similar to the case without pivot sampling that is presented in detail in (Wild et al., 2013b).

Lemma 5.3: *Conditional on the partition sizes* $\mathbf{I}$*, the number of executed Java Bytecode instructions* $T_{BC} = T_{BC}(n)$ *of the first partitioning step of* $Y_{\mathbf{t}}^w$ *— implemented as in Appendix C of (Wild et al., 2013b) — fulfills on a random permutation of size* $n > w$

$$\begin{aligned} T_{BC}(n) \;\stackrel{\mathcal{D}}{=}\; & 10n + 13I_1 + 5I_2 + 11\mathrm{HypG}(I_1 + I_2, I_3, n-k) \\ & + \mathrm{HypG}(I_1, I_1 + I_2, n-k) \,+\, O(1)\,. \end{aligned} \tag{5.6}$$

5.2.4 Distribution of Partition Sizes

There is a close relation between $\mathbf{I}$, the number of small, medium and large ordinary elements, and $\mathbf{J}$, the size of subproblems; we only have to add the sampled-out elements again before the recursive calls. So we have $\mathbf{J} = \mathbf{I} + \mathbf{t}$ and thus $\mathbb{P}(\mathbf{I} = \mathbf{i}) = \binom{i_1+t_1}{t_1}\binom{i_2+t_2}{t_2}\binom{i_3+t_3}{t_3}\big/\binom{n}{k}$. Albeit valid, this form results in nasty sums with three binomials when we try to compute expectations involving $\mathbf{I}$.

An alternative characterization of the distribution of $\mathbf{I}$ that is better suited for our needs exploits that we have i. i. d. $\mathcal{U}(0,1)$ variables. If we condition on the pivot *values*, *i.e.*, consider P and Q fixed, an ordinary element U is small, if $U \in (0, P)$, medium if $U \in (P, Q)$ and large if $U \in (Q, 1)$. The lengths $\mathbf{D} = (D_1, D_2, D_3)$ of these three intervals (see Figure 1), thus are the *probabilities* for an elements to be small, medium or large, respectively. Note that this holds *independently* of all other ordinary elements! The partition sizes $\mathbf{I}$ are then obtained as the collective outcome of $n-k$ independent drawings from this distribution, so conditional on $\mathbf{D}$, $\mathbf{I}$ is multinomially $\mathrm{Mult}(n-k, \mathbf{D})$ distributed.

With this alternative characterization, we have *decoupled* the pivot *ranks* (determined by $\mathbf{I}$) from the pivot *values*, which allows for more elegant computation of expected values (see Appendix C). This decoupling trick has (implicitly) been applied to the analysis of classic Quicksort earlier, *e. g.*, by Neininger (2001).

5.2.5 Distribution of Pivot Values

The input array is initially filled with n i. i. d. $\mathcal{U}(0,1)$ random variables from which we choose a sample $\{V_1, \dots, V_k\} \subset \{U_1, \dots, U_n\}$ of size k. The pivot values are then selected as order statistics of the sample: $P := V_{(t_1+1)}$ and $Q := V_{(t_1+t_2+2)}$ (cf. Section 3.1). In other words, $\mathbf{D}$ is the vector of *spacings* induced by the order statistics $V_{(t_1+1)}$ and $V_{(t_1+t_2+2)}$ of k i. i. d. $\mathcal{U}(0,1)$ variables $V_1, \dots, V_k$, which is known to have *Dirichlet* $\operatorname{Dir}(\mathbf{t}+1)$ distribution (Proposition B.1 in the appendix).

6 Expected Partitioning Costs

In Section 5, we characterized the full distribution of the costs of the first partitioning step. However, since those distributions are *conditional* on other random variables — namely $\mathbf{I}$, which in turn depends on $\mathbf{D}$ — we have to apply the *law of total expectation* to obtain unconditional expected values. By linearity of the expectation, it suffices to consider the summands in isolation that are collected in the following lemma. Using well-known properties of the involved distributions (collected in Appendix B for convenience), the proof becomes a standard computation; details are given in Appendix C.

Lemma 6.1: *Assume* $\mathbf{t} \in \mathbb{N}^3$, $\mathbf{D} \overset{\mathcal{D}}{=} \operatorname{Dir}(\mathbf{t}+1)$ *and* $\mathbf{I} \overset{\mathcal{D}}{=} \operatorname{Mult}(n-k, \mathbf{D})$. *Then, the following (unconditional) expectations hold:*

$$\mathbb{E}[I_j] = \frac{t_j+1}{k+1}(n-k)\,, \qquad (j=1,2,3), \tag{6.1}$$

$$\mathbb{E}\big[\mathrm{B}\big(\tfrac{I_3}{n-k}\big)\big] = \frac{t_3+1}{k+1} = \Theta(1)\,, \qquad (n\to\infty), \tag{6.2}$$

$$\mathbb{E}\big[\operatorname{HypG}(I_3, I_1, n-k)\big] = \frac{(t_1+1)(t_3+1)}{(k+1)(k+2)}(n-k-1)\,, \tag{6.3}$$

$$\mathbb{E}\big[\operatorname{HypG}(I_1+I_2, I_3, n-k)\big] = \frac{(t_1+t_2+2)(t_3+1)}{(k+1)(k+2)}(n-k-1)\,. \tag{6.4}$$

7 Solution of the Recurrence

Theorem 7.1: *Let* $\mathbb{E}[C_n]$ *be a sequence of numbers satisfying* (5.2) *for a constant* $w \geq k$ *and let the toll function* $\mathbb{E}[T_n]$ *be of the form* $\mathbb{E}[T_n] = an + O(1)$ *for a constant* a. *Then we have* $\mathbb{E}[C_n] \sim \frac{a}{\mathrm{H}(\mathbf{t})}\, n \ln n$, *where* $\mathrm{H}(\mathbf{t})$ *is given by* (4.1).

Theorem 7.1 has first been proven by Hennequin (1991, Proposition III.9) using arguments on the Cauchy-Euler differential equations that the recurrence implies for the generating function of $\mathbb{E}[C_n]$. The tool box of handy and ready-to-apply theorems has grown considerably since then. In Appendix D, we give a concise and elementary proof using the *Continuous Master Theorem* (Roura, 2001). Another possible tool that remains closer to Hennequin's original arguments is offered by Chern et al. (2002).

Theorem 4.1 now directly follows by using Lemma 6.1 on the partitioning costs from Lemma 5.1, 5.2 and 5.3 and plugging the result into Theorem 7.1.

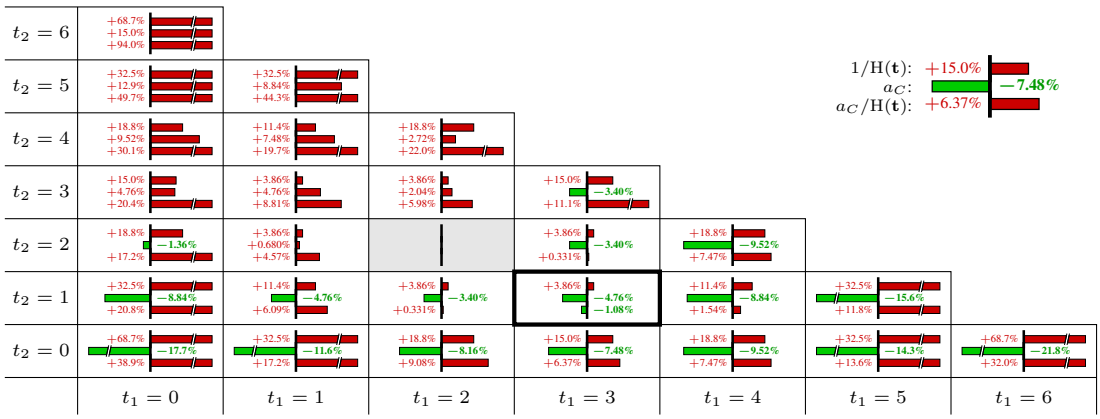

Figure 2: Inverse of discrete entropy (top), number of comparisons per partitioning step (middle) and overall comparisons (bottom) for all $\mathbf{t}$ with $k = 8$, relative to the tertiles case $\mathbf{t} = (2, 2, 2)$.

$t_1 \backslash t_2$	0	1	2	3
0	1.9956	1.8681	2.0055	2.4864
1	1.7582	**1.7043**	1.9231	
2	1.7308	1.7582		
3	1.8975			

(a) $a_C/\mathrm{H}(\mathbf{t})$

$t_1 \backslash t_2$	0	1	2	3
0	0.4907	0.4396	0.4121	**0.3926**
1	0.6319	0.5514	0.5220	
2	0.7967	0.7143		
3	1.0796			

(b) $a_S/\mathrm{H}(\mathbf{t})$

$t_1 \backslash t_2$	0	1	2	3
0	20.840	**18.791**	19.478	23.293
1	20.440	19.298	21.264	
2	22.830	22.967		
3	29.378			

(c) $a_{BC}/\mathrm{H}(\mathbf{t})$

Table 1: $\frac{a_C}{\mathrm{H}(\mathbf{t})}$, $\frac{a_S}{\mathrm{H}(\mathbf{t})}$ and $\frac{a_{BC}}{\mathrm{H}(\mathbf{t})}$ for all $\mathbf{t}$ with $k = 5$. Rows resp. columns give t_1 and t_2; t_3 is then $k - 2 - t_1 - t_2$. The symmetric choice $\mathbf{t} = (1, 1, 1)$ is shaded, the minimum is printed in bold.

8 Discussion — Asymmetries Everywhere

With Theorem 4.1, we can find the optimal sampling parameter $\mathbf{t}$ for any given sample size k. As an example, Figure 2 shows how $\mathrm{H}(\mathbf{t})$, a_C and the overall number of comparisons behave for all possible $\mathbf{t}$ with sample size $k = 8$: The discrete entropy decreases symmetrically as we move away from the center $\mathbf{t} = (2, 2, 2)$; this corresponds to the effect of less evenly distributed subproblem sizes. The individual partitioning steps, however, are cheap for *small* values of t_2 and optimal in the extreme point $\mathbf{t} = (6, 0, 0)$. For minimizing the *overall* number of comparisons — the ratio of latter two numbers — we have to find a suitable trade-off between the center and the extreme point $(6, 0, 0)$; in this case the minimal total number of comparisons is achieved with $\mathbf{t} = (3, 1, 2)$.

Apart from this trade-off between the evenness of subproblem sizes and the number of comparisons per partitioning, Table 1 shows that the optimal choices for $\mathbf{t}$ w. r. t. comparisons, swaps and Bytecodes heavily differ. The partitioning costs are, in fact, in *extreme conflict* with each other: For all $k \geq 2$, the minimal values of a_C, a_S and a_{BC} among all choices of $\mathbf{t}$ for sample size k are attained for $\mathbf{t} = (k - 2, 0, 0)$, $\mathbf{t} = (0, k - 2, 0)$ and $\mathbf{t} = (0, 0, k - 2)$, respectively. Intuitively this is so, as the strategy minimizing partitioning costs in isolation is to make the cheapest path through the partitioning loop execute as often as possible, which naturally leads to extreme choices for $\mathbf{t}$. It then depends on the actual numbers, where the total costs are minimized. It is thus not possible to minimize all cost measures at once, and the rivaling

effects described above make it hard to reason about optimal parameters merely on a qualitative level. The number of executed Bytecode instructions is certainly more closely related to actual running time than the pure number of comparisons and swaps, while it remains platform independent and deterministic.[ii] We hope that the sensitivity of the optimal sampling parameter to the chosen cost measure renews the interest in instruction-level analysis in the style of Knuth. The current focus of the literature on abstract cost measures leads to *suboptimal* choices in Yaroslavskiy's Quicksort!

It is interesting to note in this context that the implementation in Oracle's Java 7 runtime library — which uses $\mathbf{t} = (1,1,1)$ — executes asymptotically *more* Bytecodes (on random permutations) than $Y_{\mathbf{t}}^{w}$ with $\mathbf{t} = (0,1,2)$, despite using the same sample size $k = 5$. Whether this also results in a performance gain in practice, however, depends on details of the runtime environment (Wild et al., 2013c).

Continuous ranks. It is natural to ask for the optimal *relative ranks* of P and Q if we are not constrained by the discrete nature of pivot sampling. In fact, one might want to choose the sample size depending on those optimal relative ranks to find a discrete order statistic that falls close to the continuous optimum.

We can compute the optimal relative ranks by considering the limiting behavior of $Y_{\mathbf{t}}^{w}$ as $k \to \infty$. Formally, we consider the following family of algorithms: Let $(t_l^{(j)})_{j\in\mathbb{N}}$ for $l = 1,2,3$ be three sequences of non-negative integers and set $k^{(j)} := t_1^{(j)} + t_2^{(j)} + t_3^{(j)} + 2$ for every $j \in \mathbb{N}$. Assume that we have $k^{(j)} \to \infty$ and $t_l^{(j)}/k^{(j)} \to \tau_l$ with $\tau_l \in [0,1]$ for $l = 1,2,3$ as $j \to \infty$. Note that we have $\tau_1+\tau_2+\tau_3 = 1$ by definition. For each $j \in \mathbb{N}$, we can apply Theorem 4.1 for $Y_{\mathbf{t}^{(j)}}^{w}$ and then consider the limiting behavior of the total costs for $j \to \infty$.[iii] For $\mathrm{H}(\mathbf{t})$, equation (4.2) shows convergence to the entropy function $\mathrm{H}^*(\boldsymbol{\tau}) = -\sum_{l=1}^{3} \tau_l \ln(\tau_l)$ and for the numerators a_C, a_S and a_{BC}, it is easily seen that

$$a_C^{(j)} \;\to\; a_C^* \;:=\; 1 + \tau_2 + (2\tau_1+\tau_2)\tau_3\,, \tag{8.1}$$

$$a_S^{(j)} \;\to\; a_S^* \;:=\; \tau_1 + (\tau_1+\tau_2)\tau_3\,, \tag{8.2}$$

$$a_{BC}^{(j)} \;\to\; a_{BC}^* \;:=\; 10 + 13\tau_1 + 5\tau_2 + (\tau_1+\tau_2)(\tau_1+11\tau_3)\,. \tag{8.3}$$

Together, the overall number of comparisons, swaps and Bytecodes converge to $a_C^*/\mathrm{H}^*(\boldsymbol{\tau})$, $a_S^*/\mathrm{H}^*(\boldsymbol{\tau})$ resp. $a_{BC}^*/\mathrm{H}^*(\boldsymbol{\tau})$; see Figure 3 for plots. We could not find a way to compute the minima of these functions analytically. However, all three functions have isolated minima that can be approximated well by numerical methods.

The number of comparisons is minimized for $\boldsymbol{\tau}_C^* \approx (0.428846, 0.268774, 0.302380)$. For this choice, the expected number of comparisons is asymptotically $1.4931\, n\ln n$. For swaps, the minimum is not attained inside the open simplex, but for the extreme points $\boldsymbol{\tau}_S^* = (0,0,1)$ and $\boldsymbol{\tau}_S^{*\prime} = (0,1,0)$. The minimal value of the coefficient is 0, so the expected number of swaps drops to $o(n\ln n)$ for these extreme points. Of course, this is a very bad choice w. r. t. other cost measures, *e. g.*, the number of comparisons becomes quadratic, which again shows the limitations of tuning an algorithm to one of its basic operations in isolation. The minimal asymptotic number of executed Bytecodes of roughly $16.3833\, n\ln n$ is obtained for $\boldsymbol{\tau}_{BC}^* \approx (0.206772, 0.348562, 0.444666)$.

We note again that the optimal choices heavily differ depending on the employed cost measure and that the minima differ significantly from the symmetric choice $\boldsymbol{\tau} = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$.

[ii] Counting the number of executed Bytecode instructions still ignores many important effects on actual running time, *e. g.*, costs of branch mispredictions in pipelined execution, cache misses and the influence of just-in-time compilation.

[iii] Letting the sample size go to infinity implies linear overhead per partitioning for our implementation, which is, of course, not negligible any more. For the analysis here, we assume an oracle that provides us with the desired order statistic in $O(1)$.

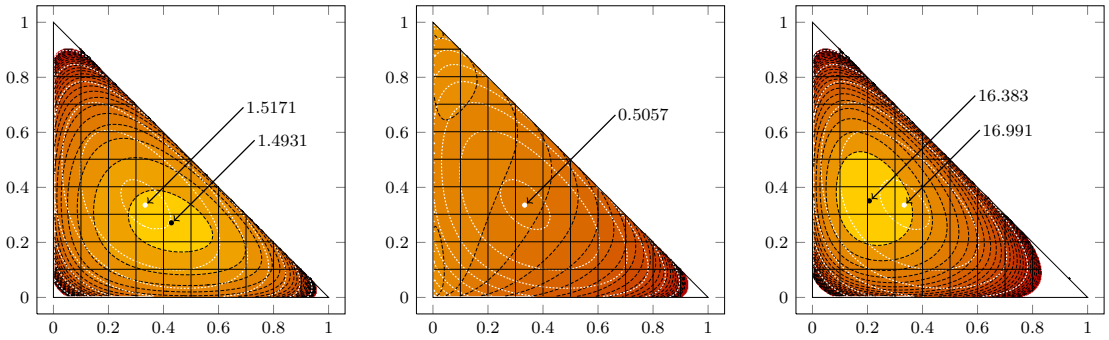

(a) $a_C^*/\mathrm{H}^*(\boldsymbol{\tau})$ (b) $a_S^*/\mathrm{H}^*(\boldsymbol{\tau})$ (c) $a_{BC}^*/\mathrm{H}^*(\boldsymbol{\tau})$

Figure 3: Contour plots for the limits of the leading term coefficient of the overall number of comparisons, swaps and executed Bytecode instructions, as functions in $\boldsymbol{\tau}$. τ_1 and τ_2 are given on x- and y-axis, respectively, which determine τ_3 as $1-\tau_1-\tau_2$. Black dots mark global minima, white dots show the center point $\tau_1=\tau_2=\tau_3=\frac{1}{3}$. (For swaps no minimum is attained in the open simplex, see main text). Black dashed lines are level lines connecting "equi-cost-ant" points, *i. e.* points of equal costs. White dotted lines mark points of equal entropy $\mathrm{H}^*(\boldsymbol{\tau})$.

9 Conclusion

In this paper, we gave the precise leading term asymptotic of the average costs of Quicksort with Yaroslavskiy's dual-pivot partitioning method and selection of pivots as arbitrary order statistics of a constant size sample. Our results confirm earlier empirical findings (Yaroslavskiy, 2010; Wild et al., 2013c) that the inherent asymmetries of the partitioning algorithm call for a systematic skew in selecting the pivots — the tuning of which requires a quantitative understanding of the delicate trade-off between partitioning costs and the distribution of subproblem sizes for recursive calls. Moreover, we have demonstrated that this tuning process is very sensitive to the choice of suitable cost measures, which firmly suggests a detailed analyses in the style of Knuth, instead of focusing on the number of comparisons and swaps only.

Future work. A natural extension of this work would be the computation of the linear term of costs, which is not negligible for moderate n. This will require a much more detailed analysis as sorting the samples and dealing with short subarrays contribute to the linear term of costs, but then allows to compute the optimal choice for w, as well. While in this paper only expected values were considered, the distributional analysis of Section 5 can be used as a starting point for analyzing the distribution of overall costs. Yaroslavskiy's partitioning can also be used in Quickselect (Wild et al., 2013a); the effects of generalized pivot sampling there are yet to be studied. Finally, other cost measures, like the number of symbol comparisons (Vallée et al., 2009; Fill and Janson, 2012), would be interesting to analyze.

References

M. Aumüller and M. Dietzfelbinger. Optimal Partitioning for Dual Pivot Quicksort. In F. V. Fomin, R. Freivalds, M. Kwiatkowska, and D. Peleg, editors, *ICALP 2013*, volume 7965 of *LNCS*, pages 33–44. Springer, 2013.

J. L. Bentley and M. D. McIlroy. Engineering a sort function. *Software: Practice and Experience*, 23(11):1249–1265, 1993.

H.-H. Chern and H.-K. Hwang. Transitional behaviors of the average cost of quicksort with median-of-$(2t+1)$. *Algorithmica*, 29(1-2):44–69, 2001.

H.-H. Chern, H.-K. Hwang, and T.-H. Tsai. An asymptotic theory for Cauchy–Euler differential equations with applications to the analysis of algorithms. *Journal of Algorithms*, 44(1):177–225, 2002.

H. A. David and H. N. Nagaraja. *Order Statistics (Wiley Series in Probability and Statistics)*. Wiley-Interscience, 3rd edition, 2003. ISBN 0-471-38926-9.

M. Durand. Asymptotic analysis of an optimized quicksort algorithm. *Information Processing Letters*, 85(2):73–77, 2003.

M. H. v. Emden. Increasing the efficiency of quicksort. *Communications of the ACM*, pages 563–567, 1970.

V. Estivill-Castro and D. Wood. A survey of adaptive sorting algorithms. *ACM Computing Surveys*, 24(4):441–476, 1992.

J. Fill and S. Janson. The number of bit comparisons used by Quicksort: an average-case analysis. *Electronic Journal of Probability*, 17:1–22, 2012.

R. L. Graham, D. E. Knuth, and O. Patashnik. *Concrete mathematics: a foundation for computer science*. Addison-Wesley, 1994. ISBN 978-0-20-155802-9.

P. Hennequin. *Analyse en moyenne d'algorithmes : tri rapide et arbres de recherche*. PhD Thesis, Ecole Politechnique, Palaiseau, 1991.

C. A. R. Hoare. Algorithm 65: Find. *Communications of the ACM*, 4(7):321–322, July 1961.

K. Kaligosi and P. Sanders. How branch mispredictions affect quicksort. In T. Erlebach and Y. Azar, editors, *ESA 2006*, pages 780–791. Springer, 2006.

S. Kushagra, A. López-Ortiz, A. Qiao, and J. I. Munro. Multi-Pivot Quicksort: Theory and Experiments. In *ALENEX 2014*, pages 47–60. SIAM, 2014.

H. M. Mahmoud. *Sorting: A distribution theory*. John Wiley & Sons, Hoboken, NJ, USA, 2000. ISBN 1-118-03288-8.

C. Martínez and S. Roura. Optimal Sampling Strategies in Quicksort and Quickselect. *SIAM Journal on Computing*, 31(3):683, 2001.

D. R. Musser. Introspective Sorting and Selection Algorithms. *Software: Practice and Experience*, 27(8):983–993, 1997.

R. Neininger. On a multivariate contraction method for random recursive structures with applications to Quicksort. *Random Structures & Algorithms*, 19(3-4):498–524, 2001.

S. Roura. Improved Master Theorems for Divide-and-Conquer Recurrences. *Journal of the ACM*, 48(2):170–205, 2001.

R. Sedgewick. *Quicksort*. PhD Thesis, Stanford University, 1975.

R. Sedgewick. The analysis of Quicksort programs. *Acta Inf.*, 7(4):327–355, 1977.

R. Sedgewick. Implementing Quicksort programs. *Communications of the ACM*, 21(10):847–857, 1978.

B. Vallée, J. Clément, J. A. Fill, and P. Flajolet. The Number of Symbol Comparisons in QuickSort and QuickSelect. In S. Albers, A. Marchetti-Spaccamela, Y. Matias, S. Nikoletseas, and W. Thomas, editors, *ICALP 2009*, volume 5555 of *LNCS*, pages 750–763. Springer, 2009.

S. Wild and M. E. Nebel. Average Case Analysis of Java 7's Dual Pivot Quicksort. In L. Epstein and P. Ferragina, editors, *ESA 2012*, volume 7501 of *LNCS*, pages 825–836. Springer, 2012.

S. Wild, M. E. Nebel, and H. Mahmoud. Analysis of Quickselect under Yaroslavskiy's Dual-Pivoting Algorithm, 2013a. URL `http://arxiv.org/abs/1306.3819`.

S. Wild, M. E. Nebel, and R. Neininger. Average Case and Distributional Analysis of Java 7's Dual Pivot Quicksort, 2013b. URL `http://arxiv.org/abs/1304.0988`.

S. Wild, M. E. Nebel, R. Reitzig, and U. Laube. Engineering Java 7's Dual Pivot Quicksort Using MaLiJAn. In P. Sanders and N. Zeh, editors, *ALENEX 2013*, pages 55–69. SIAM, 2013c.

V. Yaroslavskiy. Question on sorting. `http://mail.openjdk.java.net/pipermail/core-libs-dev/2010-July/004649.html`, 2010.

Appendix

A Detailed Pseudocode

A.1 Implementing Generalized Pivot Sampling

While extensive literature on the analysis of (single-pivot) Quicksort with pivot sampling is available, most works do not specify the pivot selection process in detail.[(iv)] The usual justification is that, in any case, we only draw pivots a *linear* number of times and from a constant size sample. So for the leading term asymptotic, the costs of pivot selection are negligible, and hence also the precise way of how selection is done is not important.

There is one caveat in the argumentation: Analyses of Quicksort usually rely on setting up a recurrence equation of expected costs that is then solved (precisely or asymptotically). This in turn requires the algorithm to *preserve* the distribution of input permutations for the subproblems subjected to recursive calls — otherwise the recurrence does not hold. Most partitioning algorithms, including the one of Yaroslavskiy, have the desirable property to preserve randomness (Wild and Nebel, 2012); but this is not sufficient! We also have to make sure that the main procedure of Quicksort does not alter the distribution of inputs for recursive calls; in connection with elaborate pivot sampling algorithms, this is harder to achieve than it might seem at first sight.

For these reasons, the authors felt the urge to include a minute discussion of how to implement the generalized pivot sampling scheme of Section 3.1 in such a way that the recurrence equation remains *precise*.[(v)] We have to address the following questions:

Which elements do we choose for the sample? In theory, a *random* sample produces the most reliable results and also protects against worst case inputs. The use of a random pivot for classic Quicksort has been considered right from its invention (Hoare, 1961) and is suggested as a general strategy to deal with biased data (Sedgewick, 1978).

However, all programming libraries known to the authors actually avoid the additional effort of drawing random samples. They use a set of deterministically selected positions of the array, instead; chosen to give reasonable results for common special cases like almost sorted arrays. For example, the positions used in Oracle's Java 7 implementation are depicted in Figure 4.

For our analysis, the input consists of i. i. d. random variables, so *all* subsets (of a certain size) have the same distribution. We might hence select the positions of sample elements such that they are convenient for our (analysis) purposes. For reasons elaborated in Section A.2 below, we have to *exclude* sampled-out elements from partitioning to keep analysis feasible, and therefore, our implementation uses the $t_1 + t_2 + 1$ leftmost and the $t_3 + 1$ rightmost elements of the array as sample, as illustrated in Figure 5. Then, partitioning can be simply restricted to the range between the two parts of the sample, namely positions $t_1 + t_2 + 2$ through $n - t_3 - 1$.

(iv) Noteworthy exceptions are Sedgewick's seminal works which give detailed code for the median-of-three strategy (Sedgewick, 1975, 1978) and Bentley and McIlroy's influential paper on engineering a practical sorting method (Bentley and McIlroy, 1993). Martínez and Roura describe a general approach of which they state that randomness is *not* preserved, but in their analysis, they "disregard the small amount of sortedness [. . .] yielding at least a good approximation" (Martínez and Roura, 2001, Section 7.2).

(v) Note that the resulting implementation has to be considered "academic": While it is well-suited for precise analysis, it will look somewhat peculiar from a practical point of view and productive use is probably not to be recommended.

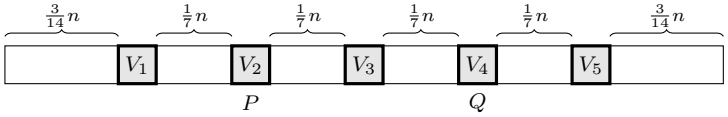

Figure 4: The five sample elements in Oracle's Java 7 implementation of Yaroslavskiy's dual-pivot Quicksort are chosen such that their distances are approximately as given above.

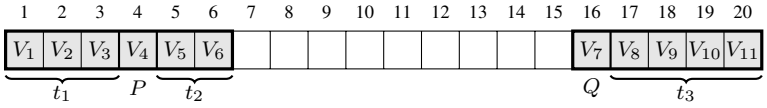

Figure 5: Location of the sample in our implementation of $Y^w_{\mathbf{t}}$ with $\mathbf{t} = (3, 2, 4)$. Only the non-shaded region $\mathtt{A}[7..15]$ is subject to partitioning.

How do we select the desired order statistics from the sample? Finding a given order statistic of a list of elements is known as the *selection problem* and can be solved by specialized algorithms like Quickselect. Even though these selection algorithms are superior by far on large lists, selecting pivots from a reasonably small sample is most efficiently done by fully sorting the whole sample with an elementary sorting method. Once the sample has been sorted, we find the pivots in $\mathtt{A}[t_1 + 1]$ and $\mathtt{A}[n - t_3]$, respectively.

We will use an Insertionsort variant for sorting samples. Note that the implementation has to "jump" across the gap between the left part and the right part of the sample. Algorithm 5 and its symmetric cousin Algorithm 6 do that by ignoring the gap for all index variables and then correct for the gap whenever the array is actually accessed.

How do we deal with sampled-out elements? As discussed in Section A.2, we exclude sampled-out elements from the partitioning range. After partitioning, we thus have to move the t_2 sampled-out elements, which actually belong between the pivots, to the middle partition. Moreover, the pivots themselves have to

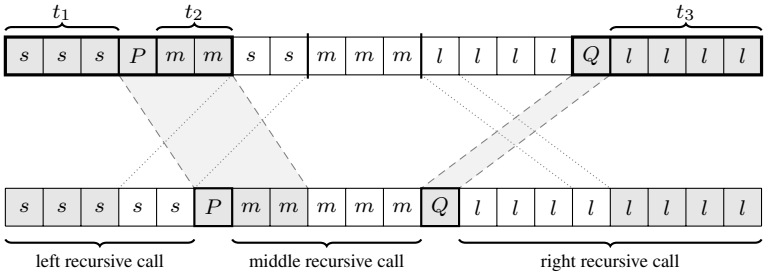
t_1
t_2
t_3
s s s P m m s s m m m l l l l Q l l l l
s s s s s P m m m m m Q l l l l l l l l
left recursive call
middle recursive call
right recursive call

Figure 6: First row: State of the array just after partitioning the ordinary elements (after line 17 of Algorithm 1). The letters indicate whether the element at this location is smaller (s), between (m) or larger (l) than the two pivots P and Q. Sample elements are shaded.
Second row: State of the array after pivots and sample parts have been moved to their partition (after line 21). The "rubber bands" indicate moved regions of the array.

be swapped in place. This process is illustrated in Figure 6 and spelled out in lines 18 – 21 of Algorithm 1. Note that the order of swaps has been chosen carefully to correctly deal with cases, where the regions to be exchanged overlap.

A.2 Randomness Preservation

For analysis, it is vital to preserve the input distribution for recursive calls, as this allows us to set up a recurrence equation for costs, which in turn underlies the precise analysis of Quicksort. While Yaroslavskiy's method (as given in Algorithm 2) preserves randomness, pivot sampling requires special care. For efficiently selecting the pivots, we *sort* the entire sample, so the sampled-out elements are far from randomly ordered; including them in partitioning would not produce randomly ordered subarrays! But there is also no need to include them in partitioning, as we already have the sample divided into the three groups of t_1 small, t_2 medium and t_3 large elements. All ordinary elements are still in random order and Yaroslavskiy's partitioning divides them into three randomly ordered subarrays.

What remains problematic is the order of elements for recursive calls. The second row in Figure 6 shows the situation after all sample elements (shaded gray) have been put into the correct subarray. As the sample was sorted, the left and middle subarrays have sorted prefixes of length t_1 resp. t_2 followed by a random permutation of the remaining elements. Similarly, the right subarray has a sorted suffix of t_3 elements. So the subarrays are *not* randomly ordered, (except for the trivial case $\mathbf{t} = 0$)! How shall we deal with this non-randomness?

The maybe surprising answer is that we can indeed *exploit* this non-randomness; not only in terms of a precise analysis, but also for efficiency: The sorted part *always* lies completely inside the *sample range* for the next partitioning phase. So our specific kind of non-randomness only affects sorting the sample (in subsequent recursive calls), but it never affects the partitioning process itself!

It seems natural that sorting should somehow be able to profit from partially sorted input, and in fact, many sorting methods are known to be *adaptive* to existing order (Estivill-Castro and Wood, 1992). For our special case of a fully sorted prefix or suffix of length $s \geq 1$ and a fully random rest, we can simply use Insertionsort where the first s iterations of the outer loop are skipped. Our Insertionsort implementations accept s as an additional parameter. What is more, we can also precisely *quantify* the savings resulting from skipping the first s iterations: Apart from per-call overhead, we save exactly what it would have costed to sort a random permutation of the length of this prefix/suffix with Insertionsort. As all prefixes/suffixes have constant lengths (independent of the length of the current subarray), precise analysis remains feasible. Thereby, we need not be afraid of non-randomness *per se*, as long as we can preserve the *same kind* of non-randomness for recursive calls and precisely analyze resulting costs.

A.3 Generalized Yaroslavskiy Quicksort

Combining the implementation of generalized pivot sampling — paying attention to the subtleties discussed in the previous sections — with Yaroslavskiy's partitioning method, we finally obtain Algorithm 1. We refer to this sorting method as *Generalized Yaroslavskiy Quicksort* with pivot sampling parameter $\mathbf{t} = (t_1, t_2, t_3)$ and Insertionsort threshold w, shortly written as $Y_{\mathbf{t}}^{w}$. We assume that $w \geq k - 1 = t_1 + t_2 + t_3 + 1$ to make sure that every partitioning step has enough elements for pivot sampling.

The last parameter of Algorithm 1 tells the current call whether it is a topmost call (`root`) or a recursive call on a left, middle or right subarray of some earlier invocation. By that, we know which part of the array is already sorted: For `root` calls, we cannot rely on anything being sorted, in `left` and `middle` calls,

we have a sorted prefix of length t_1 resp. t_2, and for a `right` call, the t_3 rightmost elements are known to be in order. The initial call then takes the form GENERALIZEDYAROSLAVSKIY $(\mathtt{A}, 1, n, \mathtt{root})$.

Algorithm 1 Yaroslavskiy's Dual-Pivot Quicksort with Generalized Pivot Sampling

GENERALIZEDYAROSLAVSKIY $(\mathtt{A}, \mathit{left}, \mathit{right}, \mathit{type})$

```
    // Assumes left ≤ right, w ≥ k − 1
    // Sorts A[left, ..., right].
 1  if right − left < w
 2      case distinction on type
 3          in case root   do INSERTIONSORTLEFT  (A, left, right, 1)
 4          in case left   do INSERTIONSORTLEFT  (A, left, right, max{t1, 1})
 5          in case middle do INSERTIONSORTLEFT  (A, left, right, max{t2, 1})
 6          in case right  do INSERTIONSORTRIGHT (A, left, right, max{t3, 1})
 7      end cases
 8  else
 9      case distinction on type    // Sort sample
10          in case root   do SAMPLESORTLEFT  (A, left, right, 1)
11          in case left   do SAMPLESORTLEFT  (A, left, right, max{t1, 1})
12          in case middle do SAMPLESORTLEFT  (A, left, right, max{t2, 1})
13          in case right  do SAMPLESORTRIGHT (A, left, right, max{t3, 1})
14      end cases
15      p := A[left + t1];   q := A[right − t3]
16      partLeft := left + t1 + t2 + 1;   partRight := right − t3 − 1
17      (ip, iq) := PARTITIONYAROSLAVSKIY (A, partLeft, partRight, p, q)

        // Swap middle part of sample and p to final place (cf. Figure 6)
18      for j := t2, ..., 0    // iterate downwards
19          Swap A[left + t1 + j] and A[ip − t2 + j]
20      end for
        // Swap q to final place.
21      Swap A[iq] and A[partRight + 1]

22      GENERALIZEDYAROSLAVSKIY (A, left,          ip − t2 − 1, left  )
23      GENERALIZEDYAROSLAVSKIY (A, ip − t2 + 1, iq − 1,        middle)
24      GENERALIZEDYAROSLAVSKIY (A, iq + 1,        right,       right )
25  end if
```

Algorithm 2 Yaroslavskiy's dual-pivot partitioning algorithm.

PARTITIONYAROSLAVSKIY $(\mathtt{A}, \mathit{left}, \mathit{right}, p, q)$

// Assumes $\mathit{left} \leq \mathit{right}$.
// Rearranges $\mathtt{A}$ s. t. with return value (i_p, i_q) holds $\begin{cases} \forall\ \mathit{left} \leq j \leq i_p, & \mathtt{A}[j] < p; \\ \forall\ i_p < j < i_q, & p \leq \mathtt{A}[j] \leq q; \\ \forall\ i_q \leq j \leq \mathit{right}, & \mathtt{A}[j] \geq q. \end{cases}$

```
1  ℓ := left;   g := right;   k := ℓ
2  while k ≤ g
3      if A[k] < p
4          Swap A[k] and A[ℓ]
5          ℓ := ℓ + 1
6      else
7          if A[k] ≥ q
8              while A[g] > q and k < g
9                  g := g − 1
10             end while
11             if A[g] ≥ p
12                 Swap A[k] and A[g]
13             else
14                 Swap A[k] and A[g];   Swap A[k] and A[ℓ]
15                 ℓ := ℓ + 1
16             end if
17             g := g − 1
18         end if
19     end if
20     k := k + 1
21 end while
22 ℓ := ℓ − 1;   g := g + 1
23 return (ℓ, g)
```

Algorithm 3 Insertionsort "from the left", exploits sorted prefixes.

INSERTIONSORTLEFT($\mathtt{A}, \mathit{left}, \mathit{right}, s$)

// Assumes $\mathit{left} \leq \mathit{right}$ and $s \leq \mathit{right} - \mathit{left} - 1$.
// Sorts $\mathtt{A}[\mathit{left}, \ldots, \mathit{right}]$, assuming that the s *leftmost* elements are already sorted.

```
1  for i = left + s, ..., right
2      j := i − 1;   v := A[i]
3      while j ≥ left ∧ v < A[j]
4          A[j + 1] := A[j];   j := j − 1
5      end while
6      A[j + 1] := v
7  end for
```

Algorithm 4 Insertionsort "from the right", exploits sorted suffixes.

INSERTIONSORTRIGHT($\mathtt{A}, \mathit{left}, \mathit{right}, s$)

// Assumes $\mathit{left} \leq \mathit{right}$ and $s \leq \mathit{right} - \mathit{left} - 1$.
// Sorts $\mathtt{A}[\mathit{left}, \ldots, \mathit{right}]$, assuming that the s *rightmost* elements are already sorted.

```
1  for i = right − s, ..., left     // iterate downwards
2      j := i + 1;   v := A[i]
3      while j ≤ right ∧ v > A[j]
4          A[j − 1] := A[j];   j := j + 1
5      end while
6      A[j − 1] := v
7  end for
```

Algorithm 5 Sorts the sample with Insertionsort "from the left"

SAMPLESORTLEFT($\mathtt{A}, \mathit{left}, \mathit{right}, s$)

// Assumes $\mathit{right} - \mathit{left} + 1 \geq k$ and $s \leq t_1 + t_2 + 1$.
// Sorts the k elements $\mathtt{A}[\mathit{left}], \ldots, \mathtt{A}[\mathit{left} + t_1 + t_2], \mathtt{A}[\mathit{right} - t_3], \ldots, \mathtt{A}[\mathit{right}]$,
// assuming that the s leftmost elements are already sorted.

// By $\mathtt{A}[\![i]\!]$, we denote the array cell $\mathtt{A}[i]$, if $i \leq \mathit{left} + t_1 + t_2$,
// and $\mathtt{A}[i + (n - k)]$ for $n = \mathit{right} - \mathit{left} + 1$, otherwise.

```
1  INSERTIONSORTLEFT(A, left, left + t1 + t2, s)
2  for i = left + t1 + t2 + 1, ..., left + k − 1
3      j := i − 1;    v := A⟦i⟧
4      while j ≥ left ∧ v < A⟦j⟧
5          A⟦j + 1⟧ := A⟦j⟧;    j := j − 1
6      end while
7      A⟦j + 1⟧ := v
8  end for
```

Algorithm 6 Sorts the sample with Insertionsort "from the right"

SAMPLESORTRIGHT($\mathtt{A}, \mathit{left}, \mathit{right}, s$)

// Assumes $\mathit{right} - \mathit{left} + 1 \geq k$ and $s \leq t_3 + 1$.
// Sorts the k elements $\mathtt{A}[\mathit{left}], \ldots, \mathtt{A}[\mathit{left} + t_1 + t_2], \mathtt{A}[\mathit{right} - t_3], \ldots, \mathtt{A}[\mathit{right}]$,
// assuming that the s rightmost elements are already sorted.

// By $\mathtt{A}[\![i]\!]$, we denote the array cell $\mathtt{A}[i]$, if $i \leq \mathit{left} + t_1 + t_2$,
// and $\mathtt{A}[i + (n - k)]$ for $n = \mathit{right} - \mathit{left} + 1$, otherwise.

```
1  INSERTIONSORTRIGHT(A, right − t3, right, s)
2  for i = left + k − t3 − 2, ..., left    // iterate downwards
3      j := i + 1;    v := A⟦i⟧
4      while j ≤ left + k ∧ v > A⟦j⟧
5          A⟦j − 1⟧ := A⟦j⟧;    j := j + 1
6      end while
7      A⟦j − 1⟧ := v
8  end for
```

B Properties of Distributions

We herein collect definitions and basic properties of the distributions used in this paper. They will be needed for computing expected values in Appendix C. We use the notation $x^{\overline{n}}$ and $x^{\underline{n}}$ of Graham et al. (1994) for rising and falling factorial powers, respectively.

B.1 Dirichlet Distribution and Beta Function

For $d \in \mathbb{N}$ let Δ_d be the standard $(d-1)$-dimensional simplex, *i. e.*,

$$\Delta_d \;:=\; \left\{ x = (x_1, \dots, x_d) \;:\; \forall i : x_i \geq 0 \;\wedge \sum_{1 \leq i \leq d} x_i = 1 \right\}. \tag{B.1}$$

Let $\alpha_1, \dots, \alpha_d > 0$ be positive reals. A random variable $\mathbf{X} \in \mathbb{R}^d$ is said to have the *Dirichlet distribution* with *shape parameter* $\boldsymbol{\alpha} := (\alpha_1, \dots, \alpha_d)$ — abbreviated as $\mathbf{X} \overset{\mathcal{D}}{=} \mathrm{Dir}(\boldsymbol{\alpha})$ — if it has a density given by

$$f_{\mathbf{X}}(x_1, \dots, x_d) \;:=\; \begin{cases} \frac{1}{\mathrm{B}(\boldsymbol{\alpha})} \cdot x_1^{\alpha_1 - 1} \cdots x_d^{\alpha_d - 1}, & \text{if } \mathbf{x} \in \Delta_d\,; \\ 0, & \text{otherwise.} \end{cases} \tag{B.2}$$

Here, $\mathrm{B}(\boldsymbol{\alpha})$ is the *d-dimensional Beta function* defined as the following Lebesgue integral:

$$\mathrm{B}(\alpha_1, \dots, \alpha_d) \;:=\; \int_{\Delta_d} x_1^{\alpha_1 - 1} \cdots x_d^{\alpha_d - 1} \, \mu(d\mathbf{x}) \,. \tag{B.3}$$

The integrand is exactly the density without the normalization constant $\frac{1}{\mathrm{B}(\alpha)}$, hence $\int f_X \, d\mu = 1$ as needed for probability distributions.

The Beta function can be written in terms of the Gamma function $\Gamma(t) = \int_0^\infty x^{t-1} e^{-x} \, dx$ as

$$\mathrm{B}(\alpha_1, \dots, \alpha_d) \;=\; \frac{\Gamma(\alpha_1) \cdots \Gamma(\alpha_d)}{\Gamma(\alpha_1 + \cdots + \alpha_d)} \,. \tag{B.4}$$

(For integral parameters $\boldsymbol{\alpha}$, a simple inductive argument and partial integration suffice to prove (B.4).)
Note that $\mathrm{Dir}(1, \dots, 1)$ corresponds to the uniform distribution over Δ_d. For integral parameters $\boldsymbol{\alpha} \in \mathbb{N}^d$, $\mathrm{Dir}(\boldsymbol{\alpha})$ is the distribution of the *spacings* or *consecutive differences* induced by appropriate order statistics of i. i. d. uniformly in $(0, 1)$ distributed random variables:

Proposition B.1 (David and Nagaraja 2003, Section 6.4): *Let $\boldsymbol{\alpha} \in \mathbb{N}^d$ be a vector of positive integers and set $k := -1 + \sum_{i=1}^d \alpha_i$. Further let $V_1, \dots, V_k$ be k random variables i. i. d. uniformly in $(0, 1)$ distributed. Denote by $V_{(1)} \leq \cdots \leq V_{(k)}$ their corresponding order statistics. We select some of the order statistics according to $\boldsymbol{\alpha}$: For $j = 1, \dots, d-1$ define $W_j := V_{(p_j)}$, where $p_j := \sum_{i=1}^j \alpha_i$. Additionally, we set $W_0 := 0$ and $W_d := 1$.*

Then, the consecutive distances (or spacings) $D_j := W_j - W_{j-1}$ for $j = 1, \dots, d$ induced by the selected order statistics $W_1, \dots, W_{d-1}$ are Dirichlet distributed with parameter $\boldsymbol{\alpha}$:

$$(D_1, \dots, D_d) \;\overset{\mathcal{D}}{=}\; \mathrm{Dir}(\alpha_1, \dots, \alpha_d) \,. \tag{B.5}$$

In the computations of Section 6, mixed moments of Dirichlet distributed variables will show up, which can be dealt with using the following general statement.

Lemma B.2: *Let* $\mathbf{X} = (X_1, \dots, X_d) \in \mathbb{R}^d$ *be a* $\mathrm{Dir}(\boldsymbol{\alpha})$ *distributed random variable with parameter* $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_d)$. *Let further* $m_1, \dots, m_d \in \mathbb{N}$ *be non-negative integers and abbreviate the sums* $A := \sum_{i=1}^{d} \alpha_i$ *and* $M := \sum_{i=1}^{d} m_i$. *Then we have*

$$\mathbb{E}\big[X_1^{m_1} \cdots X_d^{m_d}\big] \;=\; \frac{\alpha_1^{\overline{m_1}} \cdots \alpha_d^{\overline{m_d}}}{A^{\overline{M}}} \,. \tag{B.6}$$

Proof: Using $\frac{\Gamma(z+n)}{\Gamma(z)} = z^{\overline{n}}$ for all $z \in \mathbb{R}_{>0}$ and $n \in \mathbb{N}$, we compute

$$\begin{aligned}
\mathbb{E}\big[X_1^{m_1} \cdots X_d^{m_d}\big] &= \int_{\Delta_d} x_1^{m_1} \cdots x_d^{m_d} \cdot \frac{x_1^{\alpha_1 - 1} \cdots x_d^{\alpha_d - 1}}{\mathrm{B}(\boldsymbol{\alpha})} \, \mu(dx) && \text{(B.7)} \\
&= \frac{\mathrm{B}(\alpha_1 + m_1, \dots, \alpha_d + m_d)}{\mathrm{B}(\alpha_1, \dots, \alpha_d)} && \text{(B.8)} \\
&\underset{\text{(B.4)}}{=} \frac{\alpha_1^{\overline{m_1}} \cdots \alpha_d^{\overline{m_d}}}{A^{\overline{M}}} \,. && \text{(B.9)}
\end{aligned}$$

For completeness, we state here a two-dimensional Beta integral with an additional logarithmic factor that is needed in Appendix D (see also Martínez and Roura 2001, Appendix B):

$$\begin{aligned}
\mathrm{B}_{\ln}(\alpha_1, \alpha_2) &:= -\int_0^1 x^{\alpha_1 - 1} (1 - x)^{\alpha_2 - 1} \ln x \, dx && \text{(B.10)} \\
&= \mathrm{B}(\alpha_1, \alpha_2)(\mathcal{H}_{\alpha_1 + \alpha_2 - 1} - \mathcal{H}_{\alpha_1 - 1}) \,. && \text{(B.11)}
\end{aligned}$$

For integral parameters $\boldsymbol{\alpha}$, the proof is elementary: By partial integration, we can find a recurrence equation for $\mathrm{B}_{\ln}$:

$$\mathrm{B}_{\ln}(\alpha_1, \alpha_2) \;=\; \frac{1}{\alpha_1} \mathrm{B}(\alpha_1, \alpha_2) \;+\; \frac{\alpha_2 - 1}{\alpha_1} \mathrm{B}_{\ln}(\alpha_1 + 1, \alpha_2 - 1) \,. \tag{B.12}$$

Iterating this recurrence until we reach the base case $\mathrm{B}_{\ln}(a, 0) = \frac{1}{a^2}$ and using (B.4) to expand the Beta function, we obtain (B.11).

B.2 Multinomial Distribution

Let $n, d \in \mathbb{N}$ and $k_1, \ldots, k_d \in \mathbb{N}$. *Multinomial coefficients* are a multidimensional extension of binomials:

$$\binom{n}{k_1, k_2, \ldots, k_d} := \begin{cases} \dfrac{n!}{k_1! k_2! \cdots k_d!}, & \text{if } n = \sum_{i=1}^{d} k_i \ ; \\ 0, & \text{otherwise.} \end{cases} \tag{B.13}$$

Combinatorially, $\binom{n}{k_1,\ldots,k_d}$ is the number of ways to partition a set of n objects into d subsets of respective sizes $k_1, \ldots, k_d$ and thus they appear naturally in the *multinomial theorem*:

$$(x_1 + \cdots + x_d)^n = \sum_{\substack{i_1,\ldots,i_d \in \mathbb{N} \\ i_1+\cdots+i_d=n}} \binom{n}{i_1, \ldots, i_d} x_1^{i_1} \cdots x_d^{i_d} \qquad \text{for } n \in \mathbb{N} \ . \tag{B.14}$$

Let $p_1, \ldots, p_d \in [0,1]$ such that $\sum_{i=1}^{d} p_i = 1$. A random variable $\mathbf{X} \in \mathbb{N}^d$ is said to have *multinomial distribution* with parameters n and $\mathbf{p} = (p_1, \ldots, p_d)$ — written shortly as $\mathbf{X} \overset{\mathcal{D}}{=} \mathrm{Mult}(n, \mathbf{p})$ — if for any $\mathbf{i} = (i_1, \ldots, i_d) \in \mathbb{N}^d$ holds

$$\mathbb{P}(\mathbf{X} = \mathbf{i}) = \binom{n}{i_1, \ldots, i_d} p_1^{i_1} \cdots p_d^{i_d} \ . \tag{B.15}$$

We need some expected values involving multinomial variables. They can be expressed as special cases of the following mixed factorial moments.

Lemma B.3: *Let $p_1, \ldots, p_d \in [0,1]$ such that $\sum_{i=1}^{d} p_i = 1$ and consider a $\mathrm{Mult}(n, \mathbf{p})$ distributed variable $\mathbf{X} = (X_1, \ldots, X_d) \in \mathbb{N}^d$. Let further $m_1, \ldots, m_d \in \mathbb{N}$ be non-negative integers and abbreviate their sum as $M := \sum_{i=1}^{d} m_i$. Then we have*

$$\mathbb{E}\big[(X_1)^{\underline{m_1}} \cdots (X_d)^{\underline{m_d}}\big] = n^{\underline{M}} p_1^{m_1} \cdots p_d^{m_d} \ . \tag{B.16}$$

Proof: We compute

$$\begin{aligned} \mathbb{E}\big[(X_1)^{\underline{m_1}} \cdots (X_d)^{\underline{m_d}}\big] &= \sum_{\mathbf{x} \in \mathbb{N}^d} x_1^{\underline{m_1}} \cdots x_d^{\underline{m_d}} \binom{n}{x_1, \ldots, x_d} p_1^{x_1} \cdots p_d^{x_d} && \text{(B.17)} \\ &= n^{\underline{M}} p_1^{m_1} \cdots p_d^{m_d} \times && \text{(B.18)} \\ &\quad \sum_{\substack{\mathbf{x} \in \mathbb{N}^d: \\ \forall i: x_i \geq m_i}} \binom{n-M}{x_1 - m_1, \ldots, x_d - m_d} p_1^{x_1 - m_1} \cdots p_d^{x_d - m_d} && \text{(B.19)} \\ &\underset{\text{(B.14)}}{=} n^{\underline{M}} p_1^{m_1} \cdots p_d^{m_d} \ \big(\underbrace{p_1 + \cdots + p_d}_{=1}\big)^{n-M} && \text{(B.20)} \\ &= n^{\underline{M}} p_1^{m_1} \cdots p_d^{m_d} \ . && \text{(B.21)} \end{aligned}$$

C Proof of Lemma 6.1

We recall that $\mathbf{D} \stackrel{\mathcal{D}}{=} \mathrm{Dir}(\mathbf{t}+1)$ and $\mathbf{I} \stackrel{\mathcal{D}}{=} \mathrm{Mult}(n-k, \mathbf{D})$ and start with the simple ingredients: $\mathbb{E}[I_j]$ for $j = 1, 2, 3$.

$$\begin{aligned}
\mathbb{E}[I_j] &= \mathbb{E}_{\mathbf{D}}\big[\mathbb{E}[I_j \mid \mathbf{D} = \mathbf{d}]\big] && \text{(C.1)}\\
&\underset{\text{Lemma B.3}}{=} \mathbb{E}_{\mathbf{D}}\big[D_j(n-k)\big] && \text{(C.2)}\\
&\underset{\text{Lemma B.2}}{=} (n-k)\frac{t_j+1}{k+1}\,. && \text{(C.3)}
\end{aligned}$$

The term $\mathbb{E}\big[\mathrm{B}\big(\frac{I_3}{n-k}\big)\big]$ is then easily computed using (C.3):

$$\mathbb{E}\big[\mathrm{B}\big(\tfrac{I_3}{n-k}\big)\big] = \frac{\mathbb{E}[I_3]}{n-k} = \frac{t_3+1}{k+1} = \Theta(1)\,. \qquad \text{(C.4)}$$

This leaves us with the hypergeometric variables; using the well-known formula $\mathbb{E}[\mathrm{HypG}(k,r,n)] = k\frac{r}{n}$, we find

$$\begin{aligned}
\mathbb{E}\big[\mathrm{HypG}(I_1+I_2, I_3, n-k)\big] &= \mathbb{E}_{\mathbf{I}}\Big[\mathbb{E}\big[\mathrm{HypG}(i_1+i_2, i_3, n-k) \mid \mathbf{I} = \mathbf{i}\big]\Big] && \text{(C.5)}\\
&= \mathbb{E}\left[\frac{(I_1+I_2)I_3}{n-k}\right] && \text{(C.6)}\\
&= \mathbb{E}_{\mathbf{D}}\left[\frac{\mathbb{E}[I_1I_3 \mid \mathbf{D}] + \mathbb{E}[I_2I_3 \mid \mathbf{D}]}{n-k}\right] && \text{(C.7)}\\
&\underset{\text{Lemma B.3}}{=} \frac{(n-k)^{\underline{2}}\,\mathbb{E}[D_1D_3] + (n-k)^{\underline{2}}\,\mathbb{E}[D_2D_3]}{n-k} && \text{(C.8)}\\
&\underset{\text{Lemma B.2}}{=} \frac{\big((t_1+1)+(t_2+1)\big)(t_3+1)}{(k+1)^{\overline{2}}}(n-k-1)\,. && \text{(C.9)}
\end{aligned}$$

The second hypergeometric summand is obtained similarly. □

D Solution to the Recurrence

An elementary proof can be given for Theorem 7.1 using Roura's *Continuous Master Theorem* (CMT) (Roura, 2001). The CMT applies to a wide class of full-history recurrences whose coefficients can be well-approximated asymptotically by a so-called *shape function* $w : [0,1] \to \mathbb{R}$. The shape function describes the coefficients only depending on the *ratio* j/n of the subproblem size j and the current size n (not depending on n or j itself) and it smoothly continues their behavior to any real number $z \in [0,1]$. This continuous point of view also allows to compute precise asymptotics for complex discrete recurrences via fairly simple integrals.

Theorem D.1 (Martínez and Roura 2001, Theorem 18): *Let F_n be recursively defined by*

$$F_n \;=\; \begin{cases} b_n, & \text{for } 0 \le n < N; \\ t_n \;+\; \sum_{j=0}^{n-1} w_{n,j}\, F_j, & \text{for } n \ge N \end{cases} \tag{D.1}$$

where the toll function satisfies $t_n \sim K n^\alpha \log^\beta(n)$ as $n \to \infty$ for constants $K \neq 0$, $\alpha \ge 0$ and $\beta > -1$. Assume there exists a function $w : [0,1] \to \mathbb{R}$, such that

$$\sum_{j=0}^{n-1} \left| w_{n,j} \;-\; \int_{j/n}^{(j+1)/n} w(z)\, dz \right| \;=\; O(n^{-d}), \qquad (n \to \infty), \tag{D.2}$$

for a constant $d > 0$. With $H := 1 - \int_0^1 z^\alpha w(z)\, dz$, we have the following cases:

1. *If $H > 0$, then $F_n \sim \dfrac{t_n}{H}$.*
2. *If $H = 0$, then $F_n \sim \dfrac{t_n \ln n}{\tilde{H}}$ with $\tilde{H} = -(\beta+1) \int_0^1 z^\alpha \ln(z)\, w(z)\, dz$.*
3. *If $H < 0$, then $F_n \sim \Theta(n^c)$ for the unique $c \in \mathbb{R}$ with $\int_0^1 z^c w(z)\, dz = 1$.* □

The analysis of single-pivot Quicksort with pivot sampling is the application par excellence for the CMT (Martínez and Roura, 2001). We will generalize this work of Martínez and Roura to the dual pivot case.

D.1 Rewriting the Recurrence

We start from the distributional equation (5.1) by conditioning on $\mathbf{J}$. For $n > w$, this gives

$$C_n \;=\; T_n \;+\; \sum_{j=0}^{n-2} \Big(\mathbb{1}_{\{J_1=j\}} C_j \;+\; \mathbb{1}_{\{J_2=j\}} C'_j \;+\; \mathbb{1}_{\{J_3=j\}} C''_j \Big)\,. \tag{D.3}$$

Taking expectations on both sides and exploiting independence yields

$$\mathbb{E}[C_n] \;=\; \mathbb{E}[T_n] \;+\; \sum_{l=1}^{3} \sum_{j=0}^{n-2} \mathbb{E}[\mathbb{1}_{\{J_l=j\}}]\, \mathbb{E}[C_j] \tag{D.4}$$

$$\;=\; \mathbb{E}[T_n] \;+\; \sum_{j=0}^{n-2} \big(\mathbb{P}(J_1=j) + \mathbb{P}(J_2=j) + \mathbb{P}(J_3=j) \big)\, \mathbb{E}[C_j]\,, \tag{D.5}$$

which is a recurrence in the form of (D.1) with weights

$$w_{n,j} \;=\; \mathbb{P}(J_1 = j) \;+\; \mathbb{P}(J_2 = j) \;+\; \mathbb{P}(J_3 = j)\,. \tag{D.6}$$

(Note that the probabilities implicitly depend on n.)
By definition, $\mathbb{P}(J_l = j) = \mathbb{P}(I_l = j - t_l)$ for $l = 1, 2, 3$. The latter probabilities can be computed using that the marginal distribution of I_l is binomial $\mathrm{Bin}(N, D_l)$, where we abbreviate by $N := n - k$ the number of ordinary elements. It is convenient to consider $\tilde{\mathbf{D}} := (D_l, 1 - D_l)$, which is distributed like $\tilde{\mathbf{D}} \overset{\mathcal{D}}{=} \mathrm{Dir}(t_l + 1, k - t_l)$. For $i \in [0..N]$ holds

$$\begin{aligned}
\mathbb{P}(I_l = i) &= \mathbb{E}_{\mathbf{D}}\big[\mathbb{E}_{\mathbf{J}}[\mathbb{1}_{\{I_l = i\}} \mid \mathbf{D}]\big] && \text{(D.7)}\\
&= \mathbb{E}_{\mathbf{D}}\big[\tbinom{N}{i} \tilde{D}_1^i \tilde{D}_2^{N-i}\big] && \text{(D.8)}\\
&\underset{\text{Lemma B.2}}{=} \binom{N}{i} \frac{(t_l + 1)^{\overline{i}} (k - t_l)^{\overline{N-i}}}{(k+1)^{\overline{N}}}\,. && \text{(D.9)}
\end{aligned}$$

D.2 Finding a Shape Function

In general, a good guess for the shape function is $w(z) = \lim_{n\to\infty} n\, w_{n,zn}$ (Roura, 2001) and, indeed, this will work out for our weights. We start by considering the behavior for large n of the terms $\mathbb{P}(I_l = zn + r)$ for $l = 1, 2, 3$, where r does not depend on n. Assuming $zn + r \in \{0, \ldots, n\}$, we compute

$$\begin{aligned}
\mathbb{P}(I_l = zn + r) &= \binom{N}{zn + r} \frac{(t_l + 1)^{\overline{zn+r}} (k - t_l)^{\overline{(1-z)n-r}}}{(k+1)^{\overline{N}}} && \text{(D.10)}\\
&= \frac{N!}{(zn + r)!((1 - z)n - r)!} \frac{\dfrac{(zn + r + t_l)!}{t_l!} \dfrac{\big((1 - z)n - r + k - t_l - 1\big)!}{(k - t_l - 1)!}}{\dfrac{(k + N)!}{k!}} && \text{(D.11)}\\
&= \frac{k!}{t_l!(k - t_l - 1)!} \frac{(zn + r + t_l)^{\underline{t_l}} \big((1 - z)n + -r + k - t_l + 1\big)^{\underline{k-t_l-1}}}{n^{\underline{k}}}\,, && \text{(D.12)}
\end{aligned}$$

and since this is a *rational* function in n,

$$\begin{aligned}
&= (k - t_l)\binom{k}{t_l} \frac{(zn)^{t_l}((1 - z)n)^{k-t_l-1}}{n^k} \cdot \Big(1 \;+\; O(n^{-1})\Big) && \text{(D.13)}\\
&= \underbrace{(k - t_l)\binom{k}{t_l} z^{t_l}(1 - z)^{k-t_l-1}}_{=:w_l(z)} \cdot \Big(n^{-1} \;+\; O(n^{-2})\Big)\,, \qquad (n \to \infty). && \text{(D.14)}
\end{aligned}$$

Thus $n\mathbb{P}(J_l = zn) \;=\; n\mathbb{P}(I_l = zn - t_l) \;\sim\; w_l(z)$, and our candidate for the shape function is

$$w(z) \;=\; \sum_{l=1}^{3} w_l(z) \;=\; \sum_{l=1}^{3} (k - t_l)\binom{k}{t_l} z^{t_l}(1 - z)^{k-t_l-1}\,. \tag{D.15}$$

It remains to verify condition (D.2). We first note using (D.14) that

$$n w_{n,zn} \;=\; w(z) \,+\, O(n^{-1})\,. \tag{D.16}$$

Furthermore as $w(z)$ is a *polynomial* in z, its derivative exists and is finite in the compact interval $[0,1]$, so its absolute value is bounded by a constant C_w. Thus $w : [0,1] \to \mathbb{R}$ is *Lipschitz-continuous* with Lipschitz constant C_w:

$$\forall z, z' \in [0,1] \;\;:\;\; \big|w(z) - w(z')\big| \;\le\; C_w |z - z'|\,. \tag{D.17}$$

For the integral from (D.2), we then have

$$\begin{aligned}
\sum_{j=0}^{n-1} \left| w_{n,j} \,-\, \int_{j/n}^{(j+1)/n} w(z)\,dz \right| \;&=\; \sum_{j=0}^{n-1} \left| \int_{j/n}^{(j+1)/n} n w_{n,j} - w(z)\,dz \right| && \text{(D.18)}\\
&\le\; \sum_{j=0}^{n-1} \frac{1}{n} \cdot \max_{z \in \left[\frac{j}{n}, \frac{j+1}{n}\right]} \Big| n w_{n,j} - w(z) \Big| && \text{(D.19)}\\
&\underset{\text{(D.16)}}{=} \sum_{j=0}^{n-1} \frac{1}{n} \cdot \left[\max_{z \in \left[\frac{j}{n}, \frac{j+1}{n}\right]} \Big| w(j/n) - w(z) \Big| \,+\, O(n^{-1}) \right] && \text{(D.20)}\\
&\le\; O(n^{-1}) \,+\, \max_{\substack{z,z' \in [0,1]:\\ |z-z'| \le 1/n}} \big| w(z) - w(z') \big| && \text{(D.21)}\\
&\underset{\text{(D.17)}}{\le} O(n^{-1}) \,+\, C_w \frac{1}{n} && \text{(D.22)}\\
&=\; O(n^{-1})\,, && \text{(D.23)}
\end{aligned}$$

which shows that our $w(z)$ is indeed a shape function of our recurrence (with $d = 1$).

D.3 Applying the CMT

With the shape function $w(z)$ we can apply Theorem D.1 with $\alpha = 1$, $\beta = 0$ and $K = a$. It turns out that case 2 of the CMT applies:

$$\begin{aligned}
H \;&=\; 1 \,-\, \int_0^1 z\, w(z)\, dz && \text{(D.24)}\\
&=\; 1 \,-\, \sum_{l=1}^{3} \int_0^1 z\, w_l(z)\, dz && \text{(D.25)}\\
&=\; 1 \,-\, \sum_{l=1}^{3} (k - t_l) \binom{k}{t_l} \mathrm{B}(t_l + 2, k - t_l) && \text{(D.26)}\\
&\underset{\text{(B.4)}}{=} 1 \,-\, \sum_{l=1}^{3} \frac{t_l + 1}{k+1} \;=\; 0\,. && \text{(D.27)}
\end{aligned}$$

For this case, the leading term coefficient of the solution is $t_n \ln(n)/\tilde{H}$ with

$$\begin{aligned}
\tilde{H} &= -\int_0^1 z \ln(z)\, w(z)\, dz && \text{(D.28)}\\
&= \sum_{l=1}^{3} (k-t_l)\binom{k}{t_l} \mathrm{B}_{\ln}(t_l+2, k-t_l) && \text{(D.29)}\\
&\underset{\text{(B.11)}}{=} \sum_{l=1}^{3} (k-t_l)\binom{k}{t_l} \mathrm{B}(t_l+2, k-t_l)(\mathcal{H}_{k+1} - \mathcal{H}_{t_l+1}) && \text{(D.30)}\\
&= \sum_{l=1}^{3} \frac{t_l+1}{k+1}(\mathcal{H}_{k+1} - \mathcal{H}_{t_l+1})\ . && \text{(D.31)}
\end{aligned}$$

So indeed, we find $\tilde{H} = \mathrm{H}(\mathbf{t})$ as claimed in Theorem 7.1, concluding the proof.

Note that the above arguments actually *derive* — not only prove correctness of — the precise leading term asymptotics of a quite involved recurrence equation. Compared with Hennequin's original proof via generating functions, it needed much less mathematical theory.